Allocating Goods on a Graph to Eliminate Envy

Yann Chevaleyre
LAMSADE
Univ. Paris-Dauphine
France

Ulle Endriss
ILLC
University of Amsterdam
The Netherlands

Nicolas Maudet
LAMSADE
Univ. Paris-Dauphine
France

Abstract

We introduce a distributed negotiation framework for multi-agent resource allocation where interactions between agents are limited by a graph defining a negotiation topology. A group of agents may only contract a deal if that group is fully connected according to the negotiation topology. An important criterion for assessing the quality of an allocation of resources, in terms of fairness, is envy-freeness: an agent is said to envy another agent if it would prefer to swap places with that other agent. We analyse under what circumstances a sequence of deals respecting the negotiation topology may be expected to converge to a state where no agent envies any of the agents it is directly connected to. We also analyse the computational complexity of a related decision problem, namely the problem of checking whether a given negotiation state admits any deal that would both be beneficial to every agent involved and reduce envy in the agent society.

Introduction

Mechanisms for resource allocation can roughly be classified as *centralised* (when a central authority takes the responsibility to compute and allocate goods —the canonical case being auctions) or *decentralised* (with the final allocation emerging as a consequence of local, uncoordinated decisions of the agents). There are many reasons why we may want to study decentralised processes, for instance because of a lack of confidence in —or the mere absence of— a central authority. One even more compelling reason is that agents may be spatially distributed, with restricted interaction opportunities between them. In this context, the assumption of a fully connected graph (or global network) clearly vanishes. This has been recognised in various situations, for instance in the case of communication networks (discussed as an instance of a distributed constraint optimisation problem by Yokoo *et al.* (1998)), or in "real-world" economies, where "we need only shop at the corner grocery or the downtown mall to acquire goods from around the world" (Wilhite 2001).

How does this restriction on the potential interactions between agents affect the outcome of a mechanism? In general, as one may expect, it makes it more difficult to realise an *efficient* state of the system. It is then crucially important to identify the structural constraint induced by the network topology. In the context of communication networks for instance, the notion of "blocking island" has been used to abstract away from the topological structure, giving rise to a certain partitioning of the network which allows to identify bottlenecks (Frei, Faltings, & Hamdi 2005). The effect of the network topology has also been studied in economics (Bell 1996), where empirical results show how some network structures enjoy fast convergence to an efficient allocation, while others make the process much slower. In short then, while taking the topology into account is necessary in most real world applications, it usually causes difficulties.

Still, in some cases, the topology of the system can actually make things *better*. One such example is given by an interesting criterion used to assess the *fairness* of a given allocation of resources: the notion of *envy*, which we are going to detail and investigate later on in this paper. To understand the following, however, no further details are required: in the extreme case of a society of completely disconnected islands (that is, no agent can see any other agent), no agent would ever be envious of the situation of any other agent.

The case of unconnected agents being of course of little interest, we shall study in this paper how mechanisms can be designed to eliminate envy *in general*, that is with no specific assumption regarding the network structure. The rest of the paper is as follows. The next section introduces the framework of distributed resource allocation, with dealing opportunities constrained by a negotiation topology. We then precisely define what "envy" means is this context. As our approach to eliminate envy amounts to redistribute some part of the wealth generated by the trading process, we first need to study what "level" of efficiency can be guaranteed in our framework: we introduce for that purpose a notion of clique-wise efficiency, obviously weaker than global efficiency. Having established this, we prove that a simple decentralised "tax" scheme has the desired property to indeed eliminate envy. As in general it is however not possible to fully reconcile agents' self-interest and envy compensation, the solution requires some initial wealth redistribution —to compensate for a potentially unfair initial allocation of resources. In situations where this is not regarded as acceptable, it is still interesting to investigate what can be achieved within the limits of the agents' self-interest. To this end, the paper finally examines the complexity of the decision problem faced by an agent involved in such a negotiation process, namely: how difficult is it to identify a deal that is both beneficial for him, but still envy-reducing for the society?

Negotiation along Graphs

We consider the setting of a finite set of *agents* $\mathcal{A} = \{1..n\}$ negotiating over the allocation of a finite set of indivisible *resources* $\mathcal{R}$ (or *goods*). Not every agent may be able to "see" all of the other agents. A *negotiation topology* is an undirected graph $G = (\mathcal{A}, E)$, the vertices of which are the agents in $\mathcal{A}$. Two agents i and j stand in the relation E iff they can see each other. As will become clear, this means that (a) i and j may engage in negotiation and exchange resources, and (b) that i and j may envy each other (you can only interact with people, or envy people, whom you can see). Note that our visibility relation E is symmetric (the graph is undirected); this is important to be able to define negotiation along graphs in a meaningful manner, but envy along graphs could also be defined with respect to a directed graph.

An *allocation* $A : \mathcal{A} \to 2^{\mathcal{R}}$ specifies for each agent i the bundle of resources owned by i, such that each and every resource gets assigned to exactly one agent. A *deal* can be described as a pair of (distinct) allocations $\delta = (A, A')$, fixing the situation before and after the exchange. Existing work within this framework, e.g. (Sandholm 1998; Endriss *et al.* 2006; Dunne, Wooldridge, & Laurence 2005), either allows for any such deals to take place or considers structural restrictions in terms of the number of goods moving between agents. For example, a "swap deal" (Sandholm 1998) is any deal between two agents whereby exactly one good is moving from the first to the second agent, and *vice versa*. Instead, in this paper, we consider a very natural restriction induced by the negotiation topology: any deal is possible, as long as it only involves agents belonging to a common *clique* of the graph G. (A clique is a set of vertices $C \subseteq \mathcal{A}$ such that $(i, j) \in E$ for all distinct $i, j \in C$). We call deals meeting this condition *clique-deals*.

Each agent $i \in \mathcal{A}$ is equipped with a *valuation function* $v_i : 2^{\mathcal{R}} \to \mathbb{Q}$, to express their preferences over alternative bundles of resources. In the general case we do not impose any restrictions on valuation functions (such as additivity), except that they are *normalised:* $v_i(\{\}) = 0$. This assumption does not affect our results in any significant way; it merely simplifies presentation. We sometimes write $v_i(A)$ as a shorthand for $v_i(A(i))$, the value agent i assigns to the bundle it obtains in allocation A.

Deals may be coupled with monetary *side payments*. These are given in terms of a *payment function* $p : \mathcal{A} \to \mathbb{Q}$ satisfying $\sum_{i \in \mathcal{A}} p(i) = 0$. A positive value $p(i)$ means that agent i is paying; a negative value $p(i)$ means that agent i is receiving the respective amount. A deal $\delta = (A, A')$ is called *individually rational* (IR) iff there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ for all agents $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$. We assume that agents will only agree on such IR deals, *i.e.* deals that benefit everyone involved. The sum of all previous payments $p(i)$ made by agent i is given by its *payment balance* $\pi(i)$. A *negotiation state* (A, π) is a pair of an allocation and a function giving the payment balance for each agent. Finally, each agent i is equipped with a quasi-linear *utility function* u_i mapping pairs of resource bundles and past payments to a utility scale: $u_i(R, x) = v_i(R) - x$. For example, $u_i(A(i), \pi(i))$ is the utility of agent i in state (A, π), while $u_i(A(j), \pi(j))$ is the utility that i would experience if it were to swap places with j (in terms of both the bundle owned, and the sum of payments made so far).

Envy along Graphs

A central consideration when assessing alternative allocations of resources concerns *fairness* (Moulin 1988; Chevaleyre *et al.* 2006). A particular interpretation of fairness is referring to the notion that, if at all possible, we would like each agent to feel that they are at least as well of as any of the others; that is, the solution chosen should be *envy-free* (Brams & Taylor 1996). In the context of a negotiation framework admitting monetary side payments, a suitable definition of envy-freeness will have to take into account not only the allocation of resources, but also the payment balance.

We end up with the following definition: A state (A, π) is called *envy-free* iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $i, j \in \mathcal{A}$. To account for a negotiation topology, we propose a simple modification of this definition, which only considers envy amongst agents that can see each other:

Definition 1 (GEF states) *A state (A, π) is called graph-envy-free (GEF) with respect to the graph $G = (\mathcal{A}, E)$ iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $(i, j) \in E$.*

A known result states that, *in the presence of money*, envy-free states always exist (Alkan, Demange, & Gale 1991). As the restriction to a graph will only reduce the number of potentially envious agents, this existence result immediately extends to GEF states. Of course, the existence of a solution alone does not mean that we have a suitable method of finding that solution. This paper analyses to what extent the distributed negotiation approach set up in the previous section can be used to this end.

Clique-wise Efficiency

Economic efficiency of an allocation A is often measured in terms of (utilitarian) *social welfare* (Moulin 1988), defined as $sw(A) = \sum_{i \in \mathcal{A}} v_i(A)$. By a known result, a (general) deal is IR iff it increases social welfare (Endriss *et al.* 2006).

In this section we develop a notion of efficiency that takes the negotiation topology into account and relates this notion to IR negotiation when restricted to clique-deals.

Definition 2 (Clique-variants) *Let A be an allocation. Another allocation A' is called a clique-variant of A iff there exists a clique C such that $\bigcup_{i \in C} A(i) = \bigcup_{i \in C} A'(i)$ and $A(i) = A'(i)$ for all $i \notin C$.*

Observe that A and A' are clique-variants of each other iff $\delta = (A, A')$ is a clique-deal.

Definition 3 (Clique-wise efficiency) *An allocation A is called clique-wise efficient iff $sw(A) \geq sw(A')$ for every clique-variant A' of A.*

It should be noted that, in its own right, this notion of efficiency would only be of very limited interest. While our definition of envy with respect to a graph is very natural and reaching GEF states seems indeed desirable, it is questionable whether the standard notion of efficiency can be relativised with respect to a negotiation topology in a meaningful manner. Our interest in clique-wise efficiency stems from the fact that it will be helpful in in characterising conditions under which convergence to a GEF state can be guaranteed (as will become clear in the next section). But first we prove a convergence result for clique-wise efficiency:

Lemma 1 (Clique-wise efficiency) *Any sequence of IR clique-deals will eventually result in a clique-wise efficient allocation of goods.*

Proof. There can be no infinite sequence of deals (clique-deals or otherwise), because any deal strictly increases social welfare and the set of possible allocations is finite. Now let A be the terminal allocation. For the sake of contradiction, suppose that A is *not* clique-wise efficient. Then there exists an allocation A' that is a clique-variant of A such that $sw(A) < sw(A')$. But then $\delta = (A, A')$ must be IR (Endriss *et al.* 2006) as well as a clique-deal. This contradicts our assumption of A being a terminal allocation. □

Lemma 1 generalises a result by Sandholm (1998), which is equivalent to the case of a fully connected graph (in other words, the cited result applies to the case without the restrictions imposed by a negotiation topology). While Lemma 1 guarantees clique-efficient outcomes, it does not say anything about *which* clique-efficient allocation will be reached. For most graphs there will be a range of clique-efficient allocation of varying quality in terms of global efficiency (in this sense the concept is similar to that of Pareto efficiency). There is no guarantee that we will end up with the "best" clique-efficient allocation. Nevertheless, as we shall see next, clique-wise efficiency is a sufficiently strong notion to serve as a basis for negotiating envy-free states.

Convergence to GEF States

In this section we are going to prove a convergence theorem for GEF states: we show that under certain conditions on the valuation functions and for a particular choice of payment function (including a so-called "initial equitability payment"), any sequence of IR deals that respect the negotiation topology will result in a GEF state. This generalises a recent result for the case without the restrictions imposed by a negotiation topology (Chevaleyre *et al.* 2007).

We now introduce the the conditions under which our theorem applies. Firstly, all valuations v are required to be *supermodular:* $v(R_1 \cup R_2) \geq v(R_1) + v(R_2) - v(R_1 \cap R_2)$ for all $R_1, R_2 \subseteq \mathcal{R}$. That is, goods are complementaries; the valuation derived from a bundle is always at least as high as the sum of the valuations of its parts. This is a very natural class of valuations to consider, which plays an important role in several application domains. Secondly, we are going to assume that agents use a specific payment function to determine the exact payments to be made from amongst the range of IR payments that are possible in most cases. We use the *globally uniform payment function* (GUPF), which fixes payments at $p(i) = [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n$ for a given IR deal $\delta = (A, A')$. Finally, we also require agents to make so-called *initial equitability payments*. Before negotiation starts, from the initial allocation A_0, each agent is required to make a one-off payment of $\pi_0(i) = v_i(A_0) - sw(A_0)/n$. This condition does present a departure from the ideal of IR negotiation, but as argued elsewhere (Chevaleyre *et al.* 2007), and as we are also going to see later on, the concepts of envy-freeness and individual rationality are in fact incompatible: there can be no fully IR negotiation process that would guarantee convergence to a GEF state under all circumstances.

Theorem 1 (Convergence) *If all valuations are supermodular and if initial equitability payments have been made, then any sequence of IR clique-deals using the GUPF will eventually result in a GEF state.*

Proof. The proof is adapted from Chevaleyre *et al*. (2007). First observe that the use of the GUPF together with initial equitability payments ensures that we get a payment balance satisfying $\pi(i) = v_i(A) - sw(A)/n$ for every state (A, π) reached during negotiation and every agent $i \in \mathcal{A}$. For the initial state, this is just the definition of the initial equitability payments, and the GUPF payment $p(i)$ for a deal $\delta = (A, A')$ is precisely the difference between the balances associated with A' and A:

$$\begin{array}{rcl} p(i) & = & [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n \\ & = & [v_i(A') - sw(A')/n] - [v_i(A) - sw(A)/n] \\ & = & \pi'(i) - \pi(i) \end{array}$$

By Lemma 1, negotiation will eventually terminate and the final allocation A^* will be clique-wise efficient. The associated payment balance will be $\pi^*(i) = v_i(A^*) - sw(A^*)/n$. We need to show that the state (A^*, π^*) must be GEF whenever all valuations v_i are supermodular.

Let i and j be any two agents. We show that i does not envy j in state (A^*, π^*). Suppose $(i, j) \in E$ (if i cannot see j then we are done). Agents i and j form a clique. So due to the clique-wise efficiency of A^*, giving both $A^*(i)$ and $A^*(j)$ to i will not increase the sum of valuations for this clique any further:

$$v_i(A^*(i)) + v_j(A^*(j)) \quad \geq \quad v_i(A^*(i) \cup A^*(j))$$

As v_i is supermodular (and normalised) , this entails:

$$v_i(A^*(i)) + v_j(A^*(j)) \quad \geq \quad v_i(A^*(i)) + v_i(A^*(j))$$

Adding $sw(A^*)/n$ to both sides of this inequation, after some simple rearrangements, yields $u_i(A^*(i), \pi^*(i)) \geq u_i(A^*(j), \pi^*(j))$, *i.e.* agent i does indeed not envy agent j. Hence, (A^*, π^*) must be a GEF state. □

Theorem 1 means that agents can negotiate in a distributed manner, guided only by their own rational interests and limited to their "neighbourhoods" as given by the cliques of the negotiation topology, and —as long as all the side conditions are satisfied— a state that is envy-free according to all agents (whose vision is limited by the negotiation topology) will eventually emerge. In particular, agents can go ahead and negotiate any beneficial deals, without fear of getting stuck in a local optimum.

A critical point in Theorem 1 is the use of the GUPF, as this payment function does not respect the negotiation topology. However, it is easy to show that there can be *no clique-wise payment function* (a payment function giving non-zero payments only to agents belonging to a particular maximal clique within which the deal is taking place) that would allow us to achieve a convergence result for GEF states. To see this, consider the following example. Suppose there are three agents on a line (*i.e.* $E = \{(1,2),(2,1),(2,3),(3,2)\}$), the payment balance is currently 0 for all agents, and the valuation function of agent 3 is $v_3(R) = 0$ for any $R \subseteq \mathcal{R}$. Then any deal between agents 1 and 2, where the former makes a non-zero payment to the latter, will render agent 3 envious of agent 2.

A variation of this example shows that even when there exists a clique-wise payment function leading to a GEF state, the exact amount of the payments to be made may depend on agents outside the clique where the deal is taking place. For the same negotiation topology as above, let again $v_3(R) = 0$ for any $R \subseteq \mathcal{R}$, but now suppose that agent 3 has benefited from a previous deal in monetary terms, *i.e.* $\pi(3) = x$ for some $x < 0$. Then an envy-eliminating deal between agents 1 and 2 should be such that it brings the payment balance of agent 2 to at least x as well (which may or may not be possible, depending on the scenario at hand). This shows that the best possible clique-wise payment function may not be identifiable locally.

A positive point to be made about the GUPF is that the payments to non-involved agents (in particular those outside the clique where the deal is taking place) solely depend on the social surplus generated by the deal and the overall number of agents in the system. So agents do only need to be "aware" of agents they cannot "see" in so far as they need to know their overall number. This arguably corresponds well to human society: our sphere of influence may be very much restricted to a small section of society (negotiation topology), but we are still aware of some basic facts concerning society as a whole (such as the number of its members).

In fact, the GUPF appears to be the *only* payment function (meeting certain mild conditions) that would allow us to obtain a convergence result such as Theorem 1. Firstly, Lemma 1 shows that the terminal allocation A^* will be clique-wise efficient —whatever the payment function. Secondly, there are many examples where, for the final state to be GEF, the final payment balance must be defined by $\pi^*(i) = v_i(A^*) - sw(A^*)/n$, as in the proof of Theorem 1. Any scenario where all agents have the same valuation function may serve as an example. Of course, any number of payment functions could achieve the above final payment balances, as long as we can be sure that the rule applied during the very last deal is such that we get the correct values for π^*. Now, if we postulate that a *reasonable* payment function must be defined such that payments can be computed independently of whether the deal in question will be final or not, then $\pi(i) = v_i(A) - sw(A)/n$ should hold after *every* deal. This in turn, *forces* initial payments to be exactly as defined above (because the initial allocation may already be efficient and hence final), and the only possible payment function is the GUPF, as it is precisely the function we obtain when we compute the difference of the payment balances for two consecutive negotiation states.

Computational Complexity

The objectives of achieving envy-freeness and of maintaining individual rationality during negotiation are not entirely compatible; this is why Theorem 1 requires initial equitability payments. In other words, there are negotiation states such that no possible IR deal would allow us to eliminate envy (Chevaleyre *et al.* 2007). So one may ask whether a given negotiation state admits a follow-up deal that would not only be IR but that also reduces envy in the system (or, as a special case, that completely eliminates envy). This section is devoted to the analysis of the computational complexity of this problem. A precise statement of the problem, firstly, requires us to define what is to be understood by the degree of envy of a society of agents.

Define the degree to which agent i envies agent j in state (A,π) as $e(i,j) = \max\{u_i(A(j),\pi(j)) - u_i(A(i),\pi(i)), 0\}$ for $(i,j) \in E$ and $e(i,j) = 0$ for $(i,j) \notin E$. There are various options in which to combine these individual envies so as to define a notion of degree of envy of an entire society (Chevaleyre *et al.* 2007). Here, we define a *measure of envy* as any function $envy(.,.)$ mapping negotiation states to the rationals that is monotonically non-decreasing in $e(i,j)$ for any $i,j \in \mathcal{A}$, and that satisfies $envy(A,\pi) = 0$ iff (A,π) is envy-free.[1] This includes natural definitions of degree of envy, such as the *sum of all envies* $\sum_{i,j\in\mathcal{A}} e(i,j)$ and the *number of envious agents* $\#\{i \in \mathcal{A} \mid 0 < \sum_{j\in\mathcal{A}} e(i,j)\}$.

We now define the decision problem we want to analyse. For the purposes of stating the problem, let a *negotiation problem* be specified by a negotiation topology $G = (\mathcal{A}, E)$; a set of resources $\mathcal{R}$; and valuations $\langle v_1, \ldots, v_n\rangle$.

Welfare Improv. with Envy Reduction (WIER)	
Instance:	negotiation problem; state (A,π); $K \in \mathbb{Q}$; measure of envy $envy(.,.)$
Question:	Is there an alternative state (A',π') satisfying
	– $\delta = (A, A')$ is a clique-deal;
	– $v_i(A') - v_i(A) > \pi'(i) - \pi(i)$ for all $i \in \mathcal{A}$;
	– $envy(A,\pi) - envy(A',\pi') \geq K$?

That is, WIER asks whether, for a given negotiation state, there exists a clique-deal that is IR (second condition) and for which payments can be arranged such that the degree of envy reduces at least by an amount of K (third condition). Observe that, if we set $K = envy(A,\pi)$, then WIER asks whether there exists an IR clique-deal that eliminates envy entirely (independently of the measure of envy chosen).

To establish the complexity of WIER, we are going to use a reduction from another decision problem:

[1] Lastly, we also assume that any measure of envy $envy(.,.)$ is polynomial-time computable.

WELFARE IMPROVEMENT (WI)	
Instance:	negotiation problem; state (A, π)
Question:	Is there a state (A', π') such that
	– $v_i(A') - v_i(A) > \pi'(i) - \pi(i)$ for all $i \in \mathcal{A}$?

That is, WI asks whether a given state admits an IR deal (without any restriction to clique-deals). The name of the problem stems from the fact that any IR deal increases social welfare. WI has been shown to be NP-complete for a range of different ways of representing the valuation functions of agents (Dunne, Wooldridge, & Laurence 2005; Chevaleyre *et al.* 2006). As our interest here is not in preference representation languages, we state our complexity result without explicit reference to the language used; it applies to any language for which WI is NP-complete.

Theorem 2 (Complexity) *WIER is NP-complete.*

Proof. NP-membership follows from the fact that the conditions of WIER can be verified in polynomial time for any proposed solution state. We are going to show NP-hardness by reduction from WI, which is known to be NP-hard even for the case of just two agents. Given an instance of WI for two agents with valuations $\langle v_1, v_2 \rangle$, resources $\mathcal{R}$, and state (A, π), we build an instance of WIER (for three agents) as follows: The negotiation topology is $G = (\{1,2,3\}, \{(1,2),(2,1),(2,3),(3,2)\})$ and the set of resources is $\mathcal{R} \cup \{r_1^*, r_2^*, r_3^*\}$. The initial state (A^*, π^*) is defined as follows:

$$\begin{array}{ll} A^*(1) = A(1) \cup \{r_1^*\} & \pi^*(1) = \pi(1) \\ A^*(2) = A(2) \cup \{r_2^*\} & \pi^*(2) = \pi(2) \\ A^*(3) = \{r_3^*\} & \pi^*(3) = 0 \end{array}$$

The valuation functions of the first two agents are given by $v_1^* = v_1 - 3\Omega r_2^* - 3\Omega r_3^*$ and $v_2^* = v_2 - 3\Omega r_1^* - 3\Omega r_3^*$, and v_3^* is defined as follows:

$$v_3^*(R) \quad = \quad \begin{cases} \Omega & \textit{if } R = A^*(1) \textit{ or } R = A^*(2) \\ -\Omega & \textit{if } R = \{r_3^*\} \\ -2\Omega & \textit{otherwise} \end{cases}$$

Here, Ω is a constant, the value of which will be chosen big enough. Finally, we set $K = \mathit{envy}(A^*, \pi^*)$, *i.e.* we are looking for an IR clique-deal that will result in an envy-free state. This completes the description of the WIER instance. In state (A^*, π^*), the utility enjoyed by agent 3 will be $-\Omega$, and that of agents 1 and 2 will be $v_1(A) - \pi(1)$ and $v_2(A) - \pi(2)$, respectively. Now notice the following few facts:

- *In the initial state (A^*, π^*), agent 3 envies agent 2 ($e(3,2) = 2\Omega - \pi^*(2)$), and agents 1 and 2 are envy-free.*
- *No IR deal can involve agent 3*. Suppose agent 3 gets $A^*(1)$ or $A^*(2)$, then its utility increases by 2Ω. Then, either 1 or 2 would get r_3^*, which would reduce social welfare by 3Ω. Also, if 3 gets any other bundle, social welfare will decrease by an amount of Ω.
- *No IR deal can involve any of the resources $\{r_1^*, r_2^*, r_3^*\}$*. Such a deal would decrease social welfare by at least Ω.
- *A deal is IR on the WI instance iff it is IR on the WIER instance*. This is a consequence of the previous two facts.
- *Any IR deal $\delta = (A^*, A^{*\prime})$ will reduce the envy of agent 3.* Specifically, 3 can only envy 2 and the bundle held by 2 will become unattractive for agent 3 ($v_3 = -2\Omega$). That is, $e(3,2)$ will go down to $\max\{-\Omega - \pi^{*\prime}(2), 0\}$. Thus, if Ω has been chosen sufficiently big, then the envy of agent 3 will become 0 after any IR deal.
- *Agents 1 and 2 will stay envy-free after any IR deal.* This must be so because of the resources r_2^*, r_3^* (resp. r_1^*, r_3^*) they highly dislike.

The last three points imply that: if the answer to the WI problem is YES, then there is also an IR deal for the WIER instance, and any such deal is bound to bring envy down to 0. If the answer to the WI problem is NO, then there is also no IR deal for the WIER instance. □

Our proof demonstrates that the restriction to clique-deals does not affect computational complexity for this particular kind of decision problem. The same is true for the addition of the condition that the chosen deal is required to decrease envy, but this is a lot more surprising. Intuitively, finding a deal that is not just IR, but also envy-reducing, seems a lot harder than finding an IR deal. Of course, this perceived increase in complexity may both increase or decrease the complexity of checking *whether* such a deal exists (which is the question asked by WI and WIER, respectively). Theorem 2 shows that neither is the case; the problem remains NP-complete.

Related Work

Our results build on a series of papers developing an abstract negotiation framework where rational but myopic agents interact in a distributed manner (Sandholm 1998), in particular various convergence (Endriss *et al.* 2006) and complexity (Dunne, Wooldridge, & Laurence 2005) results pertaining to this framework. Specifically, our convergence result generalises a recent result (Chevaleyre *et al.* 2007) by bringing the notion of negotiation topology into the picture.

Outside of this particular line of work, fair division has mostly been addressed by economists and social scientists, e.g. (Brams & Taylor 1996); research on fair division in AI is still rare. Two exceptions are the recent work by Lipton *et al.* (2004) and Bouveret and Lang (2005). Both papers analyse the complexity of finding envy-free allocations without the possibility of monetary side payments, *i.e.* when envy-free allocation may not necessarily exist. Lipton *et al.* (2004) give a polynomial algorithm for finding an allocation with a bounded degree of envy. The measure of envy used is $\max_{i,j \in \mathcal{A}} e(i,j)$, which satisfies our criteria for measures of envy given in the previous section. Furthermore, the paper studies approximation schemes for finding allocations with minimum degree of envy. Bouveret and Lang (2005) study the computational complexity of checking whether a particular negotiation problem admits an allocation that is both envy-free (without side payments) and Pareto efficient. The paper focusses on compact encodings of the problem by means of a logic-based language for modelling the individual agent preferences.

Neither paper addresses the "procedural" aspects of finding an envy-free allocation, which is the focus of Theorem 1. In the literature on fair division in economics, the term "procedure" is usually applied to methods for finding fair allocations that are suitable for implementation by human players (Brams & Taylor 1996), which is again different from our approach, as it does not address computational issues.

Conclusion

The contributions of this paper may be summarised as follows. As a first "conceptual" contribution, we have enriched an abstract negotiation framework, previously studied by several authors (Sandholm 1998; Endriss *et al.* 2006; Dunne, Wooldridge, & Laurence 2005; Chevaleyre *et al.* 2007), with an explicit negotiation topology. A negotiation topology is an undirected graph specifying which agents may trade goods with each other. Known impossibility results (Endriss *et al.* 2006) state that (almost) any structural restriction to the negotiation protocol will mean that efficient negotiation outcomes may not be attainable anymore in all cases, and this also applies to our setting. As we have shown, this can be circumvented by using a weaker notion of efficiency (Lemma 1), but we have also argued that such a relativisation of efficiency would be only of very limited practical interest.

This is where our second conceptual contribution comes in: Concerning envy-freeness, a weakened definition that takes account of the negotiation topology *is* very much an interesting concept in practice. The straightforward relativisation of the standard definition lines up perfectly with our intuitive understanding of how envy works in communities that are not fully connected.

Finally, our technical contributions are twofold. We have firstly extended a recent convergence result to show that any sequence of individually rational deals that respect the negotiation topology will, under certain conditions, result in an allocation of goods and a payment scheme where no agent envies any of the other agents it is connected to (Theorem 1). Secondly, we have analysed the computational complexity of the problem of checking whether a given negotiation state admits a deal respecting the topology that is both individually rational and that reduces overall envy. (In this context we have also given a fairly general interpretation of what constitutes a reasonable definition of the degree of envy of a society of agents.) Our NP-completeness result (Theorem 2) illustrates that, despite the fact that convergence to an envy-free state by means of a distributed negotiation scheme is often possible, implementing a solution according to this scheme still poses an algorithmic challenge.

There are a range of interesting questions to be considered for future work. One such question concerns the complexity of checking whether a given negotiation problem admits reaching an allocation with social welfare exceeding a given threshold K, if negotiation is required to proceed by means of individually rational deals respecting the negotiation topology. If there are no restrictions on the negotiation scheme, then the problem is equivalent to the winner determination problem in combinatorial auctions, which is known to be NP-complete (Rothkopf, Pekeč, & Harstad 1998). However, with the restrictions imposed by a negotiation topology, the problem seems closer to so-called reachability properties, some of which give rise to PSPACE-complete decision problems (Dunne, Wooldridge, & Laurence 2005; Chevaleyre *et al.* 2006). In particular, the simple argument for NP-membership given in the proof of Theorem 2 does not apply here, as a "solution" would now be a sequence of deals, and we cannot be sure that the sequence in question will not have to be of exponential length.

References

Alkan, A.; Demange, G.; and Gale, D. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59(4):1023–1039.

Bell, A. M. 1996. Bilateral trading on a network: A simulation study. Presented at the Second International Conference on Computing in Economics and Finance.

Bouveret, S., and Lang, J. 2005. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. In *Proc. of IJCAI-2005*.

Brams, S., and Taylor, A. 1996. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press.

Chevaleyre, Y.; Dunne, P.; Endriss, U.; Lang, J.; Lemaître, M.; Maudet, N.; Padget, J.; Phelps, S.; Rodríguez-Aguilar, J.; and Sousa, P. 2006. Issues in multiagent resource allocation. *Informatica* 30:3–31.

Chevaleyre, Y.; Endriss, U.; Estivie, S.; and Maudet, N. 2007. Reaching envy-free states in distributed negotiation settings. In *Proc. of IJCAI-2007*.

Dunne, P.; Wooldridge, M.; and Laurence, M. 2005. The complexity of contract negotiation. *Artificial Intelligence* 164(1–2):23–46.

Endriss, U.; Maudet, N.; Sadri, F.; and Toni, F. 2006. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research* 25:315–348.

Frei, C.; Faltings, B.; and Hamdi, M. 2005. Resource allocation in communication networks using abstraction and constraint satisfaction. *IEEE Journal on Selected Areas in Communication* 23(2):304–320.

Lipton, R.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proc. of ACM-EC-2004*.

Moulin, H. 1988. *Axioms of Cooperative Decision Making*. Cambridge University Press.

Rothkopf, M.; Pekeč, A.; and Harstad, R. 1998. Computationally manageable combinational auctions. *Management Science* 44(8):1131–1147.

Sandholm, T. 1998. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*.

Wilhite, A. 2001. Bilateral trade and 'small-world' networks. *Computational Economics* 18(1):49–64.

Yokoo, M.; Durfee, E. H.; Ishida, T.; and Kuwabara, K. 1998. The distributed constraint satisfaction problem: Formalization and algorithms. *Knowledge and Data Engineering* 10(5):673–685.